CENTRE
DE RECHERCHE
SUR LES TRANSPORTS
(CRT)
CENTRE
FOR RESEARCH
ON TRANSPORTATION

UNIVERSITÉ DE MONTRÉAL
ÉCOLE DES HAUTES ÉTUDES COMMERCIALES
ÉCOLE POLYTECHNIQUE

C.P. 6128, SUCCURSALE A
MONTRÉAL (QUÉBEC)
CANADA
H3C 3J7

TÉLÉPHONE : (514) 343-7575
TÉLEX : 05-24146 BIBPOLYTEC
TÉLÉCOPIEUR (FAX) : (514) 343-7121

CRT-643

STOCHASTIC INVENTORY ROUTING: ROUTE DESIGN WITH STOCKOUTS AND ROUTE FAILURES

by

Pierre Trudeau*
Moshe Dror*[1]

* Centre de recherche sur les transports, Université de Montréal, Montreal, Québec, Canada.

[1] Industrial Engineering and Management, Ben-Gurion University of the Negev, Beer Sheva, Israel.

Université de Montréal
Centre de recherche sur les transports - Publication # 643
Août 1989

89-05955/

Dépôt légal - 3e trimestre 1989

ABSTRACT

The stochastic inventory routing problem involves the distribution of a commodity such as heating oil over a long period of time to a large set of customers. The customers maintain a local inventory of the commodity which they consume at a daily rate. Their consumption varies daily and seasonally and their exact demand is known only upon the arrival of the delivery vehicle. This paper presents a detailed analysis of this problem incorporating the stochastic nature of customers' consumptions and the possibility of route failures when the actual demand on a route exceeds the capacity of a vehicle. A number of solution procedures are compared on a large set of real life data for a period of 12 consecutive weeks. The winning strategy, though computationally more expensive, provides the best system performance and reduces (almost eliminates) the stockout phenomena.

Key-Words: inventory routing; stochastic routing; stockouts; route failures.

- - - - - - - - -

RÉSUMÉ

Le problème d'inventaire et de tournées de véhicules stochastiques implique la distribution de biens, tels que l'huile à chauffage, sur une longue période de temps et pour un grand ensemble de clients. Chaque client dispose d'un inventaire qu'il consomme à un taux quotidien qui lui est propre. Cette consommation varie avec les jours et les saisons et la demande exacte du client n'est connue qu'à l'arrivée du véhicule de livraison. Nous présentons, dans cet article, une analyse détaillée du problème qui tient compte de la nature aléatoire de la consommation des clients, mais également de la possibilité de pénurie lorsque la demande réelle sur une route excède la capacité du véhicule de livraison. Quelques approches de solution sont comparées à partir d'un large ensemble de données tirées de la réalité, et ce, sur une période de 12 semaines consécutives. La meilleure approche, bien qu'exigeant de plus longs temps de calcul, procure la meilleure performance du système en plus de réduire considérablement les pénuries chez les clients.

Mots-clés : tournées et inventaire; tournées stochastiques; clients en pénurie; véhicules en pénurie.

1. INTRODUCTION

Let's consider a number of customers in a geographic area who consume a commodity delivered to them from a centrally located depot by a distribution company. The commodity being a product such as heating oil, each customer maintains a local inventory and consumes a certain amount of that product each day. The company vehicles based at the depot are assigned each day a set of routes to resupply a subset of the customers. Each day (5 days every week) a preselected subset of customers is replenished. The customer demands are random variables which become known when a vehicle arrives at the customer's location, but only immediatly after full service is completed, as no gage is accessible to the driver. The objective is to construct daily vehicle routes which result in an 'efficient' operation over a long (yearly) time horizon taking into account the stochastic nature of customer demand parameters. This distribution planning activity is called the *stochastic inventory routing problem* (SIRP). The 'efficiency' of such a distribution operation is historically (in the industry) measured by the average number of commodity units delivered in one distribution hour.

Versions of the SIRP. are described by Anily (1986), Bell. Dalberto. Fisher. Greenfield. Jaikumar, Kedia. Mack. and Prutzman (1983), Burns. Hall, Blumenfield, and Daganzo (1985), Dror (1983), Dror, Ball, and Golden (1985/6), Dror and Levy (1986), Dror and Ball (1987), Federgruen and Zipkin (1984), Assad, Golden, Dahl, and Dror (1982, 1983) subsequently summarized in Golden, Assad, and Dahl (1984), Larson (1988), and Trudeau (1984). In a special issue edited by C.T. Hendrickson (1985), a number of papers address related inventory and routing issues in a different operational environment and perspective, and though interesting in its own right it is unrelated to the approach taken in this paper. The related research work is limited to Bell et al. (1983) and Larson (1988).

The SIRP is a large and complex logistic problem with a number of other hard combinatorial problems embedded in it. It combines a temporal component (the timing of the inventory replenishment of a customer's tank) with a spatial component (the actual routing). The timing of the replenishment is addressed by a customer selection process which in turns determines the spatial configuration for the 'local' tactical decision problem - the routing component. Those issues (timing and routing) are not independent of each other: a local spatial configuration (routing) might infer a timing of a replenishment, and a decision based on the inventory level will impact on the routing component of the problem. It is an interdependent, sequential optimization process solved every week (planning period). One measure of success is the average number of units delivered per one hour of operation (to be maximized), complemented by the number of stockouts incurred in one year (to be minimized).

2. ISSUES AND GENERAL SOLUTION APPROACH

2.1 The Inventory Issue

In a classic inventory replenishment problem such as the economic order quantity problem, the order quantity strikes a balance between the costs associated with replenishing and holding inventory. In the SIRP presented in this paper, the problem is approached from the point of view of the distribution company which is paid for the commodity upon delivery. The cost of holding inventory at customer's tank falls on the customer. For the company there is no holding cost associated with the inventory. Assuming a fixed commodity price and customer's daily consumption with no relation to the inventory level in his tank, the company would attempt to maximize the quantity replenished in order to maximize operating efficiency (and profit). This quantity is bounded by the customer's tank size. Given that the company is penalized if a customer runs out of inventory (a situation referred to as a *stockout*), and that each customer has a different demand on each day, the objective is to optimize a measure of operating efficiency while attempting to insure that no customer runs out of the commodity at any time. The impetus behind this distribution system is the importance of maintaining a sufficient supply of inventory at the customer's location in face of the cost of stockouts.

2.2 The Routing Issue

The daily replenishment problem in its deterministic version is well known by its generic name as the *vehicle routing problem* (VRP). For description and solution techniques of the VRP, the reader is referred to survey papers by Bodin et al. (1983), Christofides (1985), and more recently Laporte and Nobert (1987). The VRP *with stochastic demands* (SVRP) was examined by Stewart and Golden (1983), Dror and Trudeau (1986), Laporte, Louveaux, and Mercure (1987), and Dror, Laporte, and Trudeau, (1988). Successful solution schemes for the related multi-salesman problem are described in Gavish and Srikanth (1987). In an attempt to develop gradually the necessary terminology for the analysis of the SIRP we first restate the VRP, and then examine the SVRP.

The basic assumptions of a VRP are the following: (a) for a given set of customers their demand is known before the routes are designed, (b) each customer is supplied by exactly one vehicle, (c) the total demand on each route is less than or equal to the capacity of the vehicle assigned to that route, (d) each route begins and ends at the central depot, and in many instances there is an additional requirement that, (e) the time required to complete a route is less than or equal to some prespecified constant (usually related to the length of the driver's workday). The objective is route construction which minimizes the total distance traveled by the fleet of vehicles.

In the SVRP, the demand at each location is a random variable whose value is unknown before the routes are designed. The objective is to design a set of routes that minimizes the expected costs of the routing system. This problem is examined

in Dror, Laporte, and Trudeau (1988), and its versions are described in Dror and Trudeau (1986), Stewart and Golden (1983), and Laporte, Louveaux, and Mercure (1987). One important difference between the VRP and the SVRP is the occurrence of uncompleted vehicle routes in the case of the SVRP. An uncompleted route occurs when at some point in the route, the total actual demand of the customers up to and including that point exceeds the capacity of the vehicle. The vehicle at this point has to return to the depot to refill. Throughout the rest of the paper uncompleted route occurrence is referred to as a *route failure.*

2.3 The Combined Problem Issue

The SIRP is a long time horizon problem. Though only short period problems are solvable, they have to be formulated to reflect the attempt to minimize the delivery costs over a long term, which is not the same as minimizing the delivery costs over one or two much shorter planning periods. A short term myopic approach has the tendency to defer as many deliveries as possible to later planning periods. One of the important issues of the SIRP is therefore the 'proper' projection of the long time horizon objective into weekly (planning period) tactical planning of distribution. The planning period optimization problem actually solved becomes the problem of minimizing a temporal 'cost' of selecting the customers for service in the current planning period and assigning them to specific days, and the immediate delivery (spatial) 'cost'. Another aspect of the temporal cost is that of stockouts. If a customer is not scheduled for replenishment before the inventory in his tank runs out, a stockout occurs. Such a customer has to be replenished immediately via an emergency delivery service. The 'cost' of this emergency replenishment is much higher in comparison. A customer might even decide not to renew his contract with the distribution company if a stockout occurs.

In a large delivery problem with a long time horizon, one has to address what is called the 'front end' of the problem such as selecting a planning horizon, allocating customers to periods, and clustering rationally the vehicle service regions. Though the present study does not emphasize these issues explicitly, they are addressed implicitly in the algorithmic design of the system. We indicate in the appropriate places in this paper how these 'front end' parts of the problem are addressed. A global view of a general computerized system for planning the distribution of home heating oil is presented in Figure 1. It is a closed loop system where a distribution activity is followed by an update of the data from which new customer parameters are generated followed by customer selection and route generation procedures.

Figure 1 about here

2.4 Solution Procedure

Clearly, the SIRP is a real and important logistic problem. It serves no purpose to restate that the problem is NP-hard. In such a real life setting the analyst is concerned with selecting the appropriate methodology to solve the problem (i.e. construct efficient vehicle routes for the fleet on hand over a long time period). In this paper we present a solution approach which builds upon a methodology described in a number of papers, but instead of relying on the readers familiarity with the past work we first highlight important points of it. The SIRP implemented in those studies (in particular Dror, 1983, Dror and Levy, 1986) had the following four stages:

Stage 1: Customer selection: through the calculation of the 'best replenishment day' of each customer, construction of the set M' of customers that must be replenished during the planning period, and of the set M'' of customers that might be replenished because of spatial considerations.

Stage 2: Assign customers to days of the planning period by solving (heuristically) a generalized assignment problem.

Stage 3: Solve a vehicle routing problem for each day in the planning period using a modified Clarke and Wright algorithm.

Stage 4: Find an improved solution by executing inter-route and inter-day customer exchanges.

Throughout this procedure, a number of simplifying assumptions and measures were adopted. Vehicle routes were constructed for the set $M' \cup M''$ assuming deterministic customer demands (set at their expected value). But in fact, since the actual demand is not known ahead with certainty, there is a possibility of a route failure. This was handled through an approach adopted from Stewart and Golden (1983), which involves solving the problem with an artificially small truck capacity. Finally, in Stages 2 and 4, in order to protect against stockouts, it was never permitted to schedule a customer past its 'best replenishment day'.

2.5 Contribution of the paper

Dropping the simplifying assumptions mentioned above, the contribution of this paper is threefold :

- In stage 1, we develop an improved customer selection procedure. Previously, due to approximations, this calculation was quite sensitive to how much in advance we were computing the next replenishment day for a customer. The new procedure now uses an exact expression and shows robustness in terms of calculating the 'best replenishment day' of each customer.

- More important, in Stages 3 and 4 we develop a new algorithmic procedure which, when evaluating the expected cost of a route, takes explicitly into account the route direction, the cost of stockouts, route failures and their interrelation .without

assuming deterministic demands or relying on the artificial capacity device. As a byproduct, it is now possible to allow for replenishment of a customer beyond its 'best replenishment day'. In this algorithmic procedure, the spatial and temporal costs are accounted for simultaneously. It's important to note that part of those developments are not to be confined to an SIRP environment but are of a more general applicability.

- A simulation experiment was designed and conducted to test the above mentioned improvements. Previous results (Dror (1983)) were obtained in terms of mathematical expectations (i.e. E[number of commodity units delivered] and E[number of stockouts]). For this paper, real simulations were executed incurring real number of units delivered and real stockouts and route failures (if any). We also conducted on the same data set new computational experiments which include correlated customer demands.

The outline for the rest of this paper is as follows: In Section 3 we describe the customer selection procedure. In Section 3.1 we recall the reduction process from a long term horizon to a single period problem. The new procedure for calculating the 'best replenishment day' for each customer and the customer selection procedure appears in Section 3.2. In Section 4. we outline the route construction algorithm starting with a background section (4.1). The analysis of the expected route cost with stockouts and route failures is described in Section 4.2, and Section 4.3 generalizes some of the results to SVRP. The computational experiments are described in Section 5. Since the demands and the execution of the algorithm are simulated over a period of 12 successive weeks, we present the simulation procedure - Section 5.1. The main results are presented in Section 5.2 followed by a conclusion paragraph - Section 6.

3. CUSTOMER SELECTION

3.1 Background

As noted in the introduction, in order to maximize the company's operating efficiency, the guiding rule is to maximize the individual quantities delivered to the customers. A replenishment of inventory can only be as big as the customer's tank. What is the impact on the operating efficiency if a customer is replenished at a level less than his tank size ? Obviously, the number of replenishments in a given period would have to increase, therefore decreasing the operating efficiency. We do not know how to solve the SIRP as one problem with an annual time span. The annual distribution problem has to be reduced to a sequence of single-period problems. The key element in reducing this annual distribution problem to a sequence of single-period problems is to include penalty costs within the single-period model that reflect the long-term effect of decisions made during the single-period. This reduction process is described in Dror and Ball (1987). It is an important part of an inventory routing system and we sketch it out in some detail. It also enables to introduce elements that will be used in the description of the route construction procedure (Section 4). We start with a deterministic model and assume that all events occur in continuous time.

Consider a single customer with a fixed tank of size T and a fixed daily consumption rate μ. Denote by b the cost to refill the tank. The constraint of this model is that the customer should never stock out. If I_0 denotes the opening inventory, then the cost over an n-day period of refilling the tank when it reaches a zero level (which is the optimal policy) would be (if a fractional number of visits is allowed)

$$\frac{b(n\mu - I_0)}{T} \tag{1}$$

Now consider solving an optimization problem over an m-day planning period ($n >> m$). This optimization problem will consider all costs incurred during that period, but will also evaluate the effect of decisions made during the m-day period on the succeeding n days .

We assume that no customers require two deliveries in one planning period. There are two possible cases to consider. A customer, if not resupplied during the planning period, will either stockout or not. If I' denotes the inventory level at the beginning of the m-day period, then (a) if $I' - m\mu < 0$, the customer must be resupplied; (b) if $I' - m\mu \geq 0$, the customer need not be resupplied. In case (a) the optimal policy would dictate to resupply on the day t^* the tank becomes empty, which is $t^* = I'/\mu$ (note that we allow for t^* to be noninteger). If case (b) occurs, we will not resupply the customer. Since other considerations might influence us not to follow this policy, we evaluate the effects of deviation from it. In particular, we wish to determine c_t - the increase in the costs over the succeeding n-day period if in case (a) a customer is replenished on day t where $t < t^*$, rather than on day t^*, assuming the optimal policy is used in the succeeding n-day period, and if in case (b), a customer is replenished on day t of the m-day period rather than in a future period. In the latter case, g_t denotes the decrease in the costs over the succeeding n-day period assuming the optimal policy is used in that period.

The calculations of c_t and g_t are as follows: the closing inventory at the end of the m-day period, if the delivery is made on day t, is equal to $T - (m - t)\mu$. Using (1), the incremental future cost of delivering on day t rather than on day t^*, for $t < t^*$, is

$$b\frac{(n\mu - (T - (m - t)\mu))}{T} - b\frac{(n\mu - (T - (m - t^*)\mu))}{T} = \frac{I'}{T}b - \frac{\mu}{T}bt = c_t.$$

If in case (b) we deliver to that customer on day t, then the difference in the future costs would be

$$b\frac{(n\mu - (I' - m\mu))}{T} - b\frac{(n\mu - (T - (m - t)\mu))}{T} = b(1 - \frac{I'}{T} + \frac{t\mu}{T}) = b - c_t = g_t.$$

Note that c_t and g_t are independent of m and n. When we extend the analysis from a deterministic to a stochastic setting, we can no longer be certain to schedule a replenishment before the customer runs out of the inventory. We denote by r_l, a random variable, the amount consumed by the customer on day l. In this stochastic setting, we must decide at the beginning of the m-day period whether or not to replenish the customer. If we do not schedule a replenishment and the customer stocks out, then a penalty cost of S ($S >> b$) is incurred and the customer is automatically refilled.

The tradeoffs between an early delivery and a late delivery to a customer correspond to the higher future cost and the risk of paying a stockout penalty. If it is assumed that after the first delivery in the m-day period (whether because of a stockout or a scheduled replenishment) future deliveries will occur at the moment the tank becomes empty, then the expected cost during the next $(n + m)$ -day period associated with scheduling a delivery on day t for $1 \leq t \leq \frac{I'}{\mu}$ is given by

$$E_t = \sum_{i=1}^{t-1}(S + c_i)P_S(i) + (1 - \sum_{i=1}^{t-1} P_S(i))c_t = \sum_{i=1}^{t-1}(S + c_i - c_t)P_S(i) - c_t$$

where $P_S(t)$

$$= Pr\{\text{customer stocks out on day } t, \text{ not before, and his tank was not refilled prior to day } t\}$$

$$= Pr\{\sum_{l=1}^{t-1} r_l \leq T < \sum_{l=1}^{t} r_l\} = Pr\{R_{t-1} \leq T\} - Pr\{R_t \leq T\}$$

and $R_t = \sum_{l=1}^{t} r_l$.

The above expression of E_t is based on the idea that if a delivery is scheduled on day t, then if no stockout occurs before day t, we incur only the incremental future cost c_t. On the other hand, if a stockout occurs on day $i < t$, then we incur the stockout cost S together with c_i.

It is proven in Dror and Ball (1987), that under certain normality assumptions on the demand distribution (independence, coefficient of variation less or equal to one), the function E_t is strictly convex. In a related work (Kreimer and Dror, 1988) it is proven that this is assured for a number of other demand distributions as well.

In Figure 2 we present a typical E_t function. E_t increases at a relatively fast rate to the right of its minimal value and this rate is a function of the customer's coefficient of variation $\rho = \sigma/\mu$.

Figure 2 about here

The convexity of E_t implies a natural choice for a safety stock level, namely, the expected tank level on a day of minimal value for E_t. The ideal replenishment scheduling policy would be to service the customer always on the day E_t is minimal. This would be the case for one customer. When we move to multiple-customer model, customers interaction and finite fleet capacity might preclude such a policy.

At this point it is appropriate to indicate that a number of the 'front end' issues were answered here. Disregarding for a moment that the industry's practice is to use one week planning time horizon, we can set the planning time horizon to any number of days during which no customer will require more than one replenishment. As for the customers' selection process, it is fully outlined in Dror and Ball (1987). This process is based on the calculation of the days which correspond to the minimal value of the E_t function. For a given customer, if such a day occurs during the next planning period then he (she) is automatically selected. Other customers might also be selected based on their inventory level and relative spatial configuration.

3.2 New Procedure

We now describe the new customer selection procedure. The difference between this procedure and the one used in the past studies, though mathematically simple, has important implications in terms of the number of stockouts incurred, as later demonstrated by computational experiments. First, we reexamine the calculation of the 'best replenishment day' for a customer as presented in Dror (1983) and Dror and Ball (1987).

The evaluation of the expected cost function E_t in order to find the next best day of replenishment can be performed at any moment, starting on the day of the last replenishment (i.e. the customer's tank is full). In Dror (1983), a rule is implemented which requires the calculation of all the appropriate coefficients for each customer beginning three weeks prior to his expected stockout. The tank level for a customer at the time when the coefficients are calculated is assumed to be equal to its expected value. If we denote the day the coefficients are calculated by t_c and the expected tank level on that day by I_c then in Dror (1983), $P_S(t)$ is calculated as follows:

$$P_S(t) = Pr\{R^c_{t-1} \leq I_c\} - Pr\{R^c_t \leq I_c\} \qquad (2)$$

where $R^c_t = \sum_{l=t_c}^{t} r_l$. Expression (2) ignores the variability due to the past $t_c - 1$ days (i.e. since last replenishment). A more appropriate expression for the $P_S(t)$ is:

$$P_S(t) = Pr\{R_{t-1} \leq T < R_t | R_{t_c-1} \leq T\} \qquad (3)$$

where $R_t = \sum_{l=1}^{t} r_l$.

Calculating the t^* day using expression (2) (versus expression (3)), in case of a customer with substantial variability in his daily consumption, might have a significant effect on the determination of the 'best replenishment day' t^*. The replenishment might be scheduled later in a planning period, when in fact a stockout is likely to occur in the beginning or even prior to that planning period. We illustrate this in the following experiment:

We take from our data set one customer with tank capacity of 350 units and daily consumption rate of 2.13 units. We fix the standard deviation at three different values expressed in terms of the coefficient of variation: $\sigma/\mu = 0.1, 0.5, 0.9$. It is assumed that the r_l's are independent, normally distributed random variables.

Next, we generate the values for the E_t function and find the best replenishment day - t^*. We each time compute it using the two different expressions for $P_S(t)$. These successive calculations of the best days are presented in Figure 3 as if the calculation was performed each week, starting on day 1 with full tank.

Figure 3 here

From Figure 3 we notice that when in the calculation of the t^* day we use the expression (2) the best day moves forward in time reflecting the fact that the variability of the past is not taken into account. When the coefficient of variation is relatively small ($\sigma/\mu = 0.1$) that past variability does not affect much the successive calculations of the day t^*. But when the coefficient of variation is greater (0.5 or 0.9) the effect of ignoring the past variability is significant. When substituting the expression (3) for the expression (2), whatever coefficient of variation is used the t^* day remains constant (though in theory it could vary slightly when $Pr\{R_{t_c-1} \leq T\}$ is significantly less than one).

Once we have determined the method for calculating the best day t^* for each customer, we can proceed and select a subset of customers for replenishment one m-day period at a time. Let M denote the entire set of customers. Denote by $M' = \{i \in M : t_i^* \leq m\}$ the set of customers who must be serviced during the planning period (note that we add the index i to indicate the customer). Denote by $M'' = \{i \in M \setminus M' : t_i^*$ is no farther than four weeks away$\}$. Whether or not a customer in M'' is replenished during the planning period is a function of his spatial configuration relative to other customers in the set $M' \cup M''$ and the capacity of the fleet. Such spatial configuration is accounted for in the solution of the resulting VRP

with stochastic demands (in the past studies this is described in detail in Dror, 1983, and Dror and Levy, 1986).

Now we describe the route construction procedure.

4. ROUTE CONSTRUCTION

4.1 Background

In the studies of Dror (1983), Dror et al. (1985/6), Dror and Levy (1986), and Dror and Ball (1987), when designing the vehicle routes, the two issues of route failure and stockouts were treated using deterministic controls. First, the probability of a route failure was controlled by replacing the actual (fixed) vehicle capacity with a lower value, referred to as the *artificial capacity* (see Stewart and Golden, 1983, Larson, 1988). The vehicle routes were designed in a deterministic fashion where the customers' demands were set at their expected demand levels. The assumption was that to a great extent stochastic elements were captured in the definition of the c_{it}'s, g_{it}'s, and the t_i^*'s. The total demand (expected demand) of customers assigned to a vehicle did not exceed the value of the artificial vehicle capacity. Setting this capacity low enough would guarantee a low probability of route failure. As demonstrated in Dror and Trudeau (1986), in the case of stochastic demands it might be 'beneficial' to assign customers to a vehicle even if their total expected demand exceeds the vehicle capacity. It is simply a question of tradeoffs between the penalties for route failure versus savings. Finally, the stockout probability was controlled (Dror et al. 1985/86) by limiting the replenishment days to days no later than day t^*, where $E_{t^*} = min_t E_t$. This too has a 'cost' implication by preventing replenishments at some latter days beyond t^*.

Note that in an operational environment the dispatcher usually designs the vehicle routes at the beginning of the planning period (normally a week) for the entire period. A customer selected for replenishment during that period but stocked out on some day prior to the scheduled delivery is simply eliminated from his route without reoptimizing the route. Because of that and because the route might fail, the realization of a single vehicle route might be different from the one originally designed . Therefore, the stockout phenomena has a definite impact on the route failure phenomena and there is obviously a 'cost' implication of separating them in the stochastic analysis of the SIRP described above.

4.2 Analysis of the Expected Route Cost with Stockouts and Route Failures

Assume a route consisting of J customers that are visited in the order $1, 2, ..., J$. The calculation of the expected route cost takes place at the beginning of the planning period consisting of days 1, 2, ..., 5. In order to be general, assume that the route is planned for day k, $1 < k \leq 5$ (so that customers may stock out before day k).

In our analysis of the expected route cost, we take into account the following four cost factors:

A - the *expected stockout cost* of the customers who will not be replenished at their scheduled day (i.e. they will be replenished earlier because they stocked out).

B - the *expected routing cost* (not accounting for route failures) of the customers who will be replenished on their scheduled day (i.e. on day k) assuming that if a customer stocks out before hand, then we simply skip this customer in the routing sequence without reoptimizing the route.

C - the *expected route failure cost* for the customers who are to be replenished on their scheduled day. We assume that in case of a route failure the vehicle subsequently travels back and forth to the depot from the specific customer location where the failure occurred. Additional route failures might occur at the remaining customer locations in which case the above assumption is repeated.

D - the *incremental future cost*, i.e. the increase in the costs over the succeeding long term period if a customer is replenished prior to (or after) the optimal day, assuming the optimal policy (i.e. deterministic replenishments upon full tank consumptions) is used in the succeeding periods. In that context no future stockout is possible and the only impact on the future is via the tank level at the end of the planning period. The higher this level, the lower the number of future visits.

Notation

State l : for each customer it means that

$$\begin{cases} \text{the customer stocked out on day } l & \text{if } l \in \{1,...,k-1\}, \\ \text{the customer did not stock out before day } k & \\ \text{and will be replenished on day } k & \text{if } l = k; \end{cases}$$

$Pr_i(j)$: probability that customer i is in state j;

v_{ij} : travel cost to go from customer i to customer j; $i, j \in \{0, J+1\}$ denotes the depot;

$E_{SVRP}(l_1, l_2, ..., l_J)$: expected routing and vehicle failure cost given we are in the state $(l_1, l_2, ..., l_J)$ (i.e. customer i is in state $l_i, i = 1, ..., J$);

e : denotes a state of a route. It's a binary vector of J components with one in the i-th position (denoted by $e_i = 1$) if the i-th customer did not stock out before day k, and with zero ($e_i = 0$) indicating that he did. When appearing in a summation operation, it is more compactly expressed as $e = \sum_{i=1}^{J} e_i 2^i$;

e^i : is similar to e but restricted to the first i customers of the route, with $e_i^i = 1$;

$Pr\{e\}$: probability of route composition e;

d_{lk} : random variable denoting the quantity needed to replenish customer l on day k;

Q : vehicle capacity.

To simplify the presentation we will assume that the statistical distributions of customer demand are continuous. This eliminates the case where the vehicle becomes empty (a route failure) at the same time a customer's tank becomes full. In that case, $\sum_{l=1}^{i} d_l = jQ$ is true for a customer i and $j \leq i$. The vehicle would then return to the depot but resume its route at customer $i+1$ (if $i < J$) instead of customer i, so that v_{ii+1} would never be incurred. We also assume that the states $1, ..., k$ a customer is in are mutually exclusive (i.e. if a customer stocked out on day $j, j < l \leq k$, then he will not stockout again on day l) so that each customer can be in exactly one state. This (mild) assumption implies that

$$Pr_i(k) = 1 - \sum_{j=1}^{k-1} Pr_i(j) \quad (i = 1, 2, ..., J). \tag{4}$$

Next we assume that for different customers the events of customer states are independent of each other, i.e., $Pr\{$ customer i is in State (p) and customer j is in State $(q)\} = Pr\{$customer i is in State $(p)\} \times Pr\{$customer j is in State $(q)\}$.

At this point we can write the expected route cost expression - (**ERC**) as:

$$\mathbf{ERC} = \sum_{l_1=1}^{k} \sum_{l_2=1}^{k} ... \sum_{l_J=1}^{k} Pr_1(l_1) Pr_2(l_2) \times ... \times Pr_J(l_J) \times$$

$$\{c_{1l_1} + c_{2l_2} + ... + c_{Jl_J} + E_{SVRP}(l_1, l_2, ..., l_J)\} \tag{5}$$

where

$$c_{il_i} = \begin{cases} v_{0i} + v_{i0} + c_{il} & \text{if } l = l_i \in \{1, 2, ..., k-1\}, \\ c_{ik} & \text{if } l_i = k, \end{cases}$$

accounts for the future cost, denoted as before by c_{il}, and for the stockout cost of customer i which is $v_{0i} + v_{i0}$. Rewriting the **ERC** expression and using the fact that $\sum_{l_i=1}^{k} Pr_i(l_i) = 1$ for all $i = 1, 2, ..., J$, we obtain the following expression:

$$\mathbf{ERC} = \sum_{i=1}^{J} \sum_{l_i=1}^{k} Pr_i(l_i) c_{il_i} + \sum_{l_1=1}^{k} \sum_{l_2=1}^{k} ... \sum_{l_J=1}^{k} \prod_{i=1}^{J} Pr_i(l_i) E_{SVRP}(l_1, l_2, ..., l_J). \tag{6}$$

The first term in the expression (6) accounts for factors **A** and **D** and is relatively easy to compute. The second term, which accounts for factors **B** and **C**, requires

additional analysis. In order to do so, we separate the $E_{SVRP}(....)$ expression into two parts: the expected routing cost of the route consisting of the customers who did not stockout - (RT), and the expected penalty cost of going back and forth to customers on the route in case of route failure - (FP). Thus,

$$\begin{aligned}
&\sum_{l_1=1}^{k}\sum_{l_2=1}^{k}\cdots\sum_{l_J=1}^{k}\prod_{i=1}^{J} Pr_i(l_i)E_{SVRP}(l_1,l_2,...,l_J)\\
&\quad=\sum_{l_1=1}^{k}\sum_{l_2=1}^{k}\cdots\sum_{l_J=1}^{k}\prod_{i=1}^{J} Pr_i(l_i)RT(l_1,l_2,...,l_J)+\\
&\quad\sum_{l_1=1}^{k}\sum_{l_2=1}^{k}\cdots\sum_{l_J=1}^{k}\prod_{i=1}^{J} Pr_i(l_i)FP(l_1,l_2,...,l_J)\\
&\quad=\text{Expected cost of a regular route + Expected cost of route failures}\\
&\quad=\mathbf{R}+\mathbf{F}. \qquad (7)
\end{aligned}$$

The first term in (7) is exactly an instance of what Jaillet (1985) describes as a Probabilistic TSP (PTSP) and presents an expression which takes into account all the states at once. Denote this term by **R** and the second term by **F**. The evaluation of the term **R** is based on the following observation: the route on day k will incur a travel cost between customer i and customer j (i.e. v_{ij}), $i < j$, if and only if customer i did not stockout before day k, customer j did not stockout before day k, and customers $i+1, i+2, ..., j-1$ did stockout before day k. The probability of such event is given by the following expression:

$$Pr_i(k)(1-Pr_{i+1}(k))(1-Pr_{i+2}(k))...(1-Pr_{j-1}(k))Pr_j(k), \qquad (8)$$

therefore the overall expected regular routing cost is equal to:

$$\mathbf{R}=\sum_{i=0}^{J}\sum_{j=i+1}^{J+1} v_{ij}Pr_i(k)(1-Pr_{i+1}(k))...(1-Pr_{j-1}(k))Pr_j(k). \qquad (9)$$

To simplify the evaluation of the expected cost of route failures we assume that no more than one route failure can occur at one customer location. We can interpret it as an assumption that the tank capacity at a customer's location, for all customers, is less than the capacity of the vehicle. Since the specific composition of customers who will not stockout before day k is not known, we have to examine all possible compositions for a given set $\{1,2,...,J\}$ of customers.

Given a specific state e of a route, and following the same argument as in the calculation of the **R** expression, we incur the penalty (route failure) cost of $v_{0i}+v_{i0}$ at the i^{th} customer location ($e_i = 1$) if and only if a route failure occurs at that location. Thus,

$$\mathbf{F} = \sum_{e=1}^{2^J} \sum_{i=1}^{J} Pr\{e\} Pr\{\text{route failure occurs at } i|e\}(v_{0i} + v_{i0}) \qquad (10)$$

where $Pr\{\text{route failure occurs at } i|e\}$

$$= \begin{cases} Pr\{1^{st}\text{route failure occurs at } i \vee \\ \quad 2^{nd} \text{ route failure occurs at } i \vee \ldots \vee \\ \quad i^{th} \text{ route failure occurs at } i|e\} & \text{if } e_i = 1, \\ 0 & \text{if } e_i = 0; \end{cases}$$

$$= \begin{cases} \sum_{j=1}^{i} Pr\{j^{th}\text{route failure occurs at } i|e\} & \text{if } e_i = 1, \\ 0 & \text{if } e_i = 0. \end{cases} \qquad (11)$$

The individual probabilities of the j^{th} route failure at i can be represented simply as:

$$Pr\{\sum_{l=1}^{i-1} e_l d_{lk} \leq jQ < \sum_{l=1}^{i} e_l d_{lk}\}$$

where $d_{lk} = \min\{T_l, \sum_{t=t_0+1}^{k} r_{lt}\}$ - the demand on day k given that t_0 denotes the last replenishment day. Thus, the expression for $\mathbf{F}$ is:

$$\sum_{e=1}^{2^J} \sum_{i=1}^{J} \sum_{j=1}^{i} Pr\{e\}(Pr\{\sum_{l=1}^{i-1} e_l d_{lk} \leq jQ < \sum_{l=1}^{i} e_l d_{lk}\})(v_{0i} + v_{i0}). \qquad (12)$$

We can simplify the expression (12) by observing that the $Pr\{\text{route failure occurs at } i\}$ is independent of the presence of customers $i+1, i+2, \ldots, J$ in the route. Thus,

$$\mathbf{F} = \sum_{i=1}^{J} \sum_{e^i=1}^{2^{i-1}} Pr\{e^i\} Pr\{\text{route failure occurs at } i|e^i\}(v_{0i} + v_{i0})$$

$$= \sum_{i=1}^{J} \sum_{e^i=1}^{2^{i-1}} Pr\{e^i\}(\sum_{j=1}^{i} Pr\{\sum_{l=1}^{i-1} e_l^i d_{lk} \leq jQ < \sum_{l=1}^{i} e_l^i d_{lk}\})(v_{0i} + v_{i0}) \qquad (13)$$

where $\sum_{e^i=1}^{2^{i-1}}$ goes through all the states involving the first i customers for which $e_i^i = 1$. Note that $\sum_{i=1}^{J} \sum_{e^i=1}^{2^{i-1}}$ describes all the different $2^J - 1$ states except the one where every customer stocked out before the replenishment day.

At this point, using (6),(7),(9) and (13), the analytic expression for the expected route cost - **ERC** can be stated in a computable form:

$$\mathbf{ERC} = \sum_{i=1}^{J}\sum_{l_i=1}^{k} Pr_i(l_i)c_{il_i} + \sum_{i=0}^{J}\sum_{j=i+1}^{J+1} v_{ij}Pr_i(k)(1-Pr_{i+1}(k))...(1-Pr_{j-1}(k))Pr_j(k)$$

$$+\sum_{i=1}^{J}\sum_{e^i=1}^{2^{i-1}} Pr\{e^i\}(\sum_{j=1}^{i} Pr\{\sum_{l=1}^{i-1} e_l^i d_{lk} \le jQ < \sum_{l=1}^{i} e_l^i d_{lk}\})(v_{0i}+v_{i0}). \qquad (14)$$

To account for the probabilistic nature of customers'demands more explicitly in the construction of the vehicle routes (Stages 3 and 4) new expected route cost and expected node-interchange savings are established using (14). Let's denote the expected route cost with customer i originally on it by ERC(i), and by ERC($i;j$) the expected route cost of the new route visiting customer i, given i and j swapped routes. In Stage 4, when evaluating the swapping of positions of customers i and j on their respective route, the expected savings ES($i;j$) is now given by :

$$ES(i;j) = ERC'(i) + ERC'(j) - ERC(i;j) - ERC'(j;i) \qquad (15)$$

In this expected savings term, when interchanging the positions of two customers we account for all the four factors (**A**, **B**, **C**, and **D**) which constituted our expected route cost, i.e., expected stockout cost, expected route costs, expected route failure costs, and the future incremental costs. The modified Clarke and Wright savings are evaluated accordingly in Stage 3.

At this point we address a 'front end' issue of the inventory routing problem: the clustering of the vehicle service regions. By clustering of service regions it is usually understood some nonovelapping geographic partition where each region has a somewhat 'regular' shape and a vehicle route visits customers in one region. Here, this regional partition is generated daily through the vehicle routing algorithm without aiming explicitly for a 'regular' partition. When examining the vehicle routes constructed, the clustering properties of the routing algorithm demonstrate themselves. The individual routes had almost no crossovers and served quite 'regular' geographic regions.

4.3 Generalization to SVRP

The developments made in section 4.2 to obtain an expression that takes into account the **B** and **C** factors can be useful in a more general SVRP context. Suppose, for an instant, that one wants to deliver a set $\{1,2,...,J\}$ of customers, where each of them has a stochastic demand with a known distribution. Furthermore, suppose that one seeks an expression for the expected routing cost associated with visiting them in a given sequence, say $1,2,...,J$. Suppose also that the stockout phenomenon is not an issue here as each customer has to be visited.

Stewart and Golden (1983) propose a number of SVRP models, either of the chance-constrained or of the stochastic programming with recourse form. None of them though takes explicitly into account the possibility for the vehicle of returning to the depot for replenishment in case of a route failure. The only penalty is through the expected number of units exceeding vehicle capacity. In our notation, the expected cost expression would be

$$\sum_{i=0}^{J} v_{ii+1} + \lambda E[\sum_{l=1}^{J} d_l - Q] \tag{16}$$

where

d_l : the r.v. denoting the demand of customer l;

λ : the penalty per unit of demand in excess of vehicle capacity.

Note that (16) does not reflect either the situation where the tour stops as soon as all the demand have been satisfied or the first route failure occurs because then there always is a probability that the cost $\sum_{i=0}^{J} v_{ii-1}$ will not be fully incurred, a situation not accounted for by (16).

In Dror and Trudeau (1986), the authors present an algorithm (a heuristic) which constructs vehicle routes with the objective of minimizing the expected routing costs while taking into account route failures and route directions. It is assumed that the back and forth travel cost will be incurred for each of the customers on the route at and after the location of the first route failure.

In section 4.2, when a route failure occurs the empty vehicle returns to the depot for replenishment and resumes its route at the location of the failure, without any reoptimization of the sequence of the remaining customers. Contrary to Dror and Trudeau (1986), additional route failures might occur, each time with the same back and forth trip occuring. The expected routing cost expression is then :

$$\sum_{i=0}^{J} v_{ii+1} + \sum_{i=1}^{J} \sum_{j=1}^{i} Pr[\sum_{l=1}^{i-1} d_l \leq jQ < \sum_{l=1}^{i} d_l](v_{0i} + v_{i0}).$$

The first term represents the regular tour cost that will be incurred whatever the demands are, whereas the second term is the expected penalty cost due to the potential route failures.

As in section 4.2, this expression is valid only for continuous demand distributions, but is easily generalized.

5. COMPUTATIONAL EXPERIMENTS

Early work on the distribution of home heating oil (of which this paper is a continuation of a sort) started in the Fall of 1981, and was presented in Assad et al. (1982, 1983), (another reference is Golden et al., 1984). For those studies, a large, representative real world data base was collected (from one district area in Pennsylvania) and it served as a test case in a number of subsequent studies. Since this is a representative data base which contains information for more than 2,000 customers and is used to estimate their consumption parameters for each month in a 12 month period, it is a much more complete test sample than a deterministic (static) test problem for which heuristic solutions are designed and empirically tested. Approximately 20 customers are generally replenished each day in this setting.

Based on the historical observations about the consumption behaviour of the customers in the data base, consumption parameters such as mean daily consumption rates, individual seasonality factors and consumption variability are calculated in Assad et al. (1983), using an adaptation of an exponential smoothing technique. The rest of the problem parameters such as the locations, tank sizes, vehicle capacity, speed, etc. were also obtained from the actual problem setting. The same procedure for estimating the consumption parameters is used in this study.

There are some differences in the computational experiments presented here and the ones in the work by Dror, and Dror et al., (1983, 1985/86, 1986, 1987). In the latter, the customers' "actual" demands were estimated using their expected values. Here the actual demands of customers on a route are obtained by simulating their demands utilizing their estimated distribution parameters. Thus, actual route failures and stockouts can be observed and accounted for. At all stages of the route design, the actual vehicle capacity is used and not a lower - 'artificial' vehicle capacity.

In this paper, the computational comparison is between three solution approaches. The first solution approach is essentially a rerun of the algorithmic scheme presented in Dror (1983) with some modifications. We refer to this solution approach as **I**. The algorithmic details of this approach are described in Dror et al. (1985/86), Dror and Levy (1986), and Dror and Ball (1987). The modifications to that solution scheme involve the algorithmic component of customer assignment to days. In Dror (1983), a linear programming based heuristic was used to solve a generalized assignment problem close to optimality. Here a simple sorting type assignment rule is implemented.

The second solution approach, denoted by **II**, is closely related to the first approach. They differ essentially in the way the 'best' replenishment day for a customer is calculated. In the solution approach **II** the best day is computed using conditional probabilities (i.e., expression (3) is used to calculate the stockout probabilities instead of (2) which is used in **I**). The purpose for comparing **I** and **II** is to evaluate the impact that an exact calculation of t^* day and customer selection have on the operating efficiency of the routing system. We are particularly interested to see whether or not the new procedure will help to reduce the number of stockouts since the old

procedure was tending to generate t^* days passed the true ones. The third solution procedure, denoted by **III**, adopts the calculation of customers' 'best replenishment days' from the procedure **II**, and implements a route redesign scheme which fully accounts for the stochastic nature of the problem as described in Section 4.2.

5.1 Description of the Simulation

Since the computational results for all three algorithmic approaches **I**, **II**, and **III**, were obtained by simulating the actual customer demands realized upon vehicle's arrival at a location, we first describe the simulation.

The inventory routing experiment is conducted for a time length of 12 consecutive weekly periods during the winter season. A database consisting of 2077 customers is used. Each customer i is characterized by the following parameters:

T_i - the tank capacity,

μ_i - the average daily consumption.

σ_i^2 - the variance of the daily consumption.

s_{ij} - the seasonality factor for month j (j = 1,2, ..., 12),

$(corx_i, cory_i)$ - the location of the customer,

b_i - the estimated cost (time) of including the customer in a route (see for example Dror and Ball (1987) for a more detailed interpretation of this cost),

ta_i^0 - the average tank level at the beginning of the simulation.

During the simulation, other parameters are calculated:

ta_i - the average tank level at the beginning of the current day,

tr_i - the actual tank level at the beginning of the current day,

nd_i - the number of days since the last replenishment as of the beginning of the current day.

The initial values ta_i and tr_i are obtained by using respectively the average daily consumption rates and seasonality factors since the last replenishment, and by simulating the daily consumptions from the appropriate normal distributions. Note that it was possible for a customer to have a negative tr_i value which corresponded to a stockout before the simulation even started. We indeed encountered 6 such cases. For these 6 customers we imposed the following initial conditions: $ta_i = tr_i = T_i$ and $nd_i = 0$.

In the simulation phase the following assumptions hold, for each week :

• The routing operations for a week are planned just before the start of day 1 of the planning period.

• On a day that a replenishment occurs, the consumption during that day is added in that refill as if the replenishment occured at the end of that day.

• The inventory routing operation is only from Monday till Friday (inclusive). No delivery can occur on Saturday or Sunday.

• A customer who stocked out before the replenishment day is simply removed from his route and no reoptimization of that route is performed.

• The number of vehicles is two and their capacity is 2210 gallons (22,100 units) apiece. A vehicle can run more than one route on any given day. In the actual setting the company also served a small number of industrial customers who required much more frequent deliveries. These industrial customers were served by a dedicated vehicle with considerable slack for emergency service.

• When a customer stocks out, a special delivery is assigned to replenish his tank. This delivery takes place at the beginning of the day of the stockout. Each stockout generates a back and forth trip from the depot to the customer.

• When a route failure occurs, the vehicle makes a back and forth trip to the depot.

• In addition to the time it takes to refill the customer's tank - at a rate of 25 gallons per minute, a penalty of 5 minutes is added for the paper work involved.

• The vehicle travels at a constant speed of 30 mph and the distances are adjusted euclidian distances which correspond to the 'true' travel distances (as in Assad et al., 1982,1983, and Dror, 1983).

For each planning period we obtain a set of routes assigned to specific days using the specific solution approach; then the 'real' daily consumption of each customer is generated, the routes are executed and the stockouts and route failures (if any) are registered and accounted for. Note that in the results presented in Table I below, the number of units delivered per hour of operation include the 5 minute delay associated with visiting a customer and the time used in case of stockouts and route failures.

In Table I, we present the results obtained for the three approaches by generating vehicle routes for the stochastic inventory routing problem described in this paper. The results are obtained for 12 consecutive weeks (from the 1st of January).

5.2 The Main Results

At the time of data collection (early 1981), in the real life setting the average

Week	Approach I		Approach II		Approach III	
	Units/hour	cpu time sec.	Units/hour	cpu time sec.	Units/hour	cpu time sec.
1	5464	675	5685	658	5925	5607
2	6698	181	5932	297	6367	2350
3	5787	388	5478	427	5663	4788
4	5394	467	5728	377	5933	4874
5	5171	301	5913	378	5803	5799
6	6738	205	6266	260	6047	3933
7	5890	341	6185	404	6648	2570
8	6326	292	5697	333	6690	2436
9	5870	346	5687	344	6413	1872
10	5700	322	5567	395	5420	13065
11	5758	292	5628	376	5507	8017
12	5203	234	5245	355	5962	5100
average	5820	337	5765	384	6043	5034
	No. of stockouts	No. of route failures	No. of stockouts	No. of route failures	No. of stockouts	No. of route failures
total	7	3	0	1	1	1

Table I : Summary of the computational results for a period of 12 consecutive weeks.

operating efficiency of the heating oil distribution system was below 4200 units/hour with over 55 stockouts for 12 weeks of operation.

The average number of units/hour for each of the solution procedures presented in Table I is obtained by adding the total number of units delivered in the period of 12 weeks and dividing that number by the total number of distribution hours. It is not an average of the 12 weekly entries in that table. The computer used for generating the results for approach I, II, and III was a CYBER 173 using PASCAL compiler. It is a relatively slow main frame computational facility and one should relate to the reported cpu times accordingly (a division factor of 2 or 3 would be appropriate for a more modern facility).

By examining the results presented in Table I we conclude that approach II reduces (eliminates in this case) the number of stockouts and the number of route failures. The average efficiency of the operation decreases slightly (by less than 1 %). By using t^* days passed the true ones, approach I is incurring more stockouts and

thus more penalties and delays that diminish the efficiency of the system, but at the same time delivers larger quantities to each customer, contributing to increase the efficiency of the system. The average cpu time for approach II is somewhat higher than in approach I (384 seconds versus 337 seconds). This average cpu time increases by about one order of magnitude when processing the approach III. This solution procedure is indeed computationally quite costly (even though we introduced a number of thresholds in order not to consider the less significant states), as was expected from the route construction algorithm since it accounts more fully for the stochastic nature of this SIRP problem. The payoff is in terms of the system's efficiency (almost 4 % improvement over approach I and almost 5 % improvement over approach II). In addition, we notice only one route failure and one stockout.

For the computational results presented in Table I, we have assumed uncorrelated customer demands. Since the problem examined in this paper focuses on home heating oil delivery, it is important to examine the system's performance if a week of unseasonably cold (warm) weather simultaneously increases (decreases) the daily consumption rates of all the customers by 15 %. We choose week no. 5 (beginning of February) as the week with the unseasonable climatic change and rerun all three approaches with the actual (simulated) daily consumption increased (decreased) by 15 % for all customers. The results of this computational tests are reported in Table II (in average number of units/hour for the entire 12 weeks, stockouts, and route failures).

Consumption	Approach I	Approach II	Approach III
+ 15 %	5713 (u/h)	5697 (u/h)	5806 (u/h)
- 15 %	5779 (u/h)	5722 (u/h)	5998 (u/h)
Stockouts			
+ 15 %	25	5	25
- 15 %	7	0	1
Route Failures			
+ 15 %	3	1	0
- 15%	3	1	1

Table II : The results for correlated demands.

From the results presented in Table II, we can conclude that as expected, the number of stockouts increases when there is a sudden decrease in the temperature

(cold weather). The most affected is approach III which attempts to account more fully for the customers consumption parameters in its replenishment decisions and allows for deliveries passed t^* days. The fact that the route failures are unaffected by the weather changes is quite interesting. Indeed, when the temperature increases, the number of route failures in approach III even drops from one to zero. This result illustrates very well the interrelation between the stockout and route failure phenomenons : as the consumption and hence demands of customers go up in cold weather, the number of route failures first increases . When the increase in demand is high enough, then the number of stockouts starts to increase which in turn means that, because of the way the system handles stockouts (immediate emergency replenishment and removal from the scheduled route) fewer of the regularly scheduled customers need to be served on the non-emergency routes. The vehicle then has an ample supply of inventory for this smaller number of customers and so route failures are still rare, as shown by Table II.

For the regular computational tests presented in Table I we also computed some vehicle statistics in terms of average, maximum, and minimum number of vehicle operating minutes per day. Those numbers are respectively 163, 586, and 39 for approach I, 156, 555, and 32 for approach II, and 163, 397, and 52 for approach III.

6. CONCLUSION

In this study, we examined a very important logistic problem for constructing delivery routes for stochastic inventory routing. This paper examines the stochastic inventory routing problem in a greater detail than the previously reported studies, permitting a better understanding of the interdependent factors which impact on the efficiency of the delivery system. Many of the simplifying assumptions which were made in past studies of the SIRP are no longer made here. The impact of such assumptions (in customer selection procedure, artificial capacity, estimates of the expected number of stockouts and route failures) is evaluated empirically within the context of a real life SIRP. Since at the present (and foreseeable future) no computationally tractable exact solution approach for such a large complex problem is in sight , no estimates are available on how far the proposed solutions for the SIRP are from an optimal one. All we can measure is how much better the proposed solutions are from the industry standards and from previously proposed solutions. The industry standards in terms of units per hour and the number of stockouts are far below those which can be obtained with systems proposed in this paper. As far as past studies are concerned, this paper puts in clearer perspective the question of obtaining a more efficient delivery system for SIRP. The question answered here is what kind of operating efficiency can the managers aim for and at what computational effort. The main results, which are summarized in Table I, illustrate that improved understanding of the problem leads to improved system efficiency, including control of stockouts and route failures. Whether such an improved management of the delivery of home heating oil is of greater or lesser significance to a distribution company is not ours to answer.

Acknowledgment: This paper benefited greatly from the comments made by the referees on its previous versions and from comments of Michel Gendreau. We are very thankful.

REFERENCES

(1) Anily, S., "Integrating Inventory Control and Transportation Planning", Ph.D. Thesis, Graduate School of Business, Columbia University (1986).

(2) Assad, A., Golden, B., Dahl, R., and Dror, M., "Design of an Inventory/Routing System for Large Propane-Distribution Firm", **Proceedings of the 1982 Southeast TIMS Conference**, (Edited by Carl Gooding), pp. 315-320, (1982).

(3) Assad, A., Golden, B., Dahl, R., and Dror, M., "Evaluating the Effectiveness of an Integrated System for Fuel Delivery", **Proceedings of the 1983 Southeast TIMS Conference**, (Edited by John Eatman), pp. 153-160, (1983).

(4) Bell, W., Dalberto, L., Fisher, M., Greenfield, A., Jaikumar, R., Kedia, P., Mack, R., and Prutzman, P., "Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer", **Interfaces**, 13,6, pp. 1-4, (1983).

(5) Bodin, L., Golden, B., Assad, A., and Ball, M., "Routing and scheduling of vehicles and crews: The state of the art", **Computers & Operations Research**, 10, 2, pp.1-62 (1983).

(6) Christofides, N., "Vehicle routing", **The Traveling Salesman Problem**. (Edited by E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan, and D.B. Shmoys), John Wiley and Sons Ltd., pp. 431-444, (1985).

(7) Burns, L.D., Hall, D.E., Blumenfeld, D.E., and Daganzo, C.F., "Distribution strategies that minimize transportation and inventory costs", **Operations Research**, 33, pp. 469-490, (1985).

(8) Dror, M., "The Inventory Routing Problem", Working Paper MS/S, College of Business and Management, University of Maryland at College Park, 1983.

(9) Dror, M., and Ball, M., "Inventory routing: Reduction from an annual to short period problem", **Naval Research Logistics**, 34, pp. 891-905, (1987).

(10) Dror, M., and Levy, L., "A vehicle routing improvement algorithm comparison of a 'greedy' and a 'matching' implementation for inventory routing", **Computers & Operations Research**, 13, 1, pp. 33-45, (1986).

(11) Dror, M., and Trudeau, P., "Stochastic vehicle routing with modified savings algorithm", **European J. of Operational Research**, 23, pp. 228-235, (1986).

(12) Dror, M., Ball, M., and Golden, B., "A computational comparison of algorithms for the inventory routing problem", **Annals of Operations Research**, 4, pp. 3-23, (1985/6).

(13) Dror, M., Laporte, G., and Trudeau, P., "Vehicle routing with stochastic demands: models and issues", **Transportation Science**, August 1989.

(14) Federgruen, A., and Zipkin, P., "A combined vehicle routing and inventory allocation problem", **Operations Research**, 32, 1019,(1984).

(15) Gavish, B., and Srikanth, K., "An optimal solution method for large-scale multiple traveling salesman problems", **Operations Research**, 34, 5, pp. 698-717, (1987).

(16) Golden, B., Assad, A., and Dahl, R., "Analysis of a large-scale vehicle routing problem with an inventory component", **Large Scale Systems**, 7, pp. 181-190, (1984).

(17) Hendrickson, C.T. (ed.), "Transportation systems and logistics", Special issue, **Transportation Research, B**, Vol. 19B, 5, October 1985.

(18) Jaillet, P., "Probabilistic Traveling Salesman Problem", Ph. D. Thesis, Massachusetts Institute of Technology (1985).

(19) Kreimer, J., and Dror, M., "The monotonicity of threshold detection probability in some stochastic accumulation processess", (to appear in **Computers & Operations Research**), (1989).

(20) Laporte, G., Nobert, Y., "Exact algorithms for the vehicle routing problem", in S. Martello, G. Laporte, M. Minoux, and C. Ribeiro, eds. *Surveys in Combinatorial Optimization, Annals of Discrete Mathematics* 31, North-Holland Publishing, pp. 147-184, Amsterdam (1987).

(21) Laporte G., Louveaux, F., and Mercure, H., "Models and exact solutions for a class of stochastic location-routing problems", **European Journal of Operational Research**, 39, pp. 71-78, (1989).

(22) Larson, R.C., "Transportation sludge to the 106 mile site: an inventory/routing algorithm for fleet sizing and logistics system design", **Transportation Science**, Vol. 22, 3, pp. 186-198, (1988).

(23) Stewart, W.R., Jr., and Golden, B., "Stochastic vehicle routing: A comprehensive approach", **European J. of Operational Research**, 14, pp. 371-285, (1983).

(24) Trudeau, P., "Reapprovisionnement Automatique en Combustible", Universite de Montreal, Centre de recherche sur les transports - Publication No. 347, April 1984.

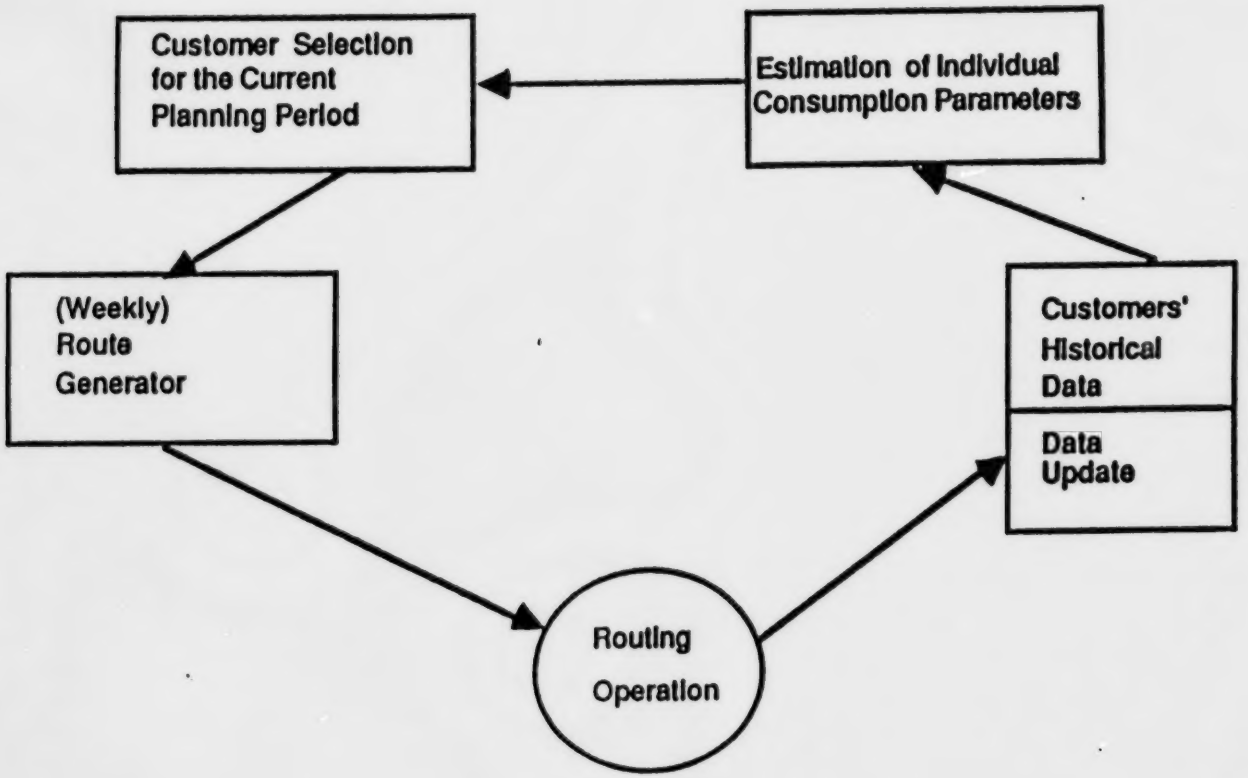

Figure 1: Operational diagram for distribution of home heating oil.

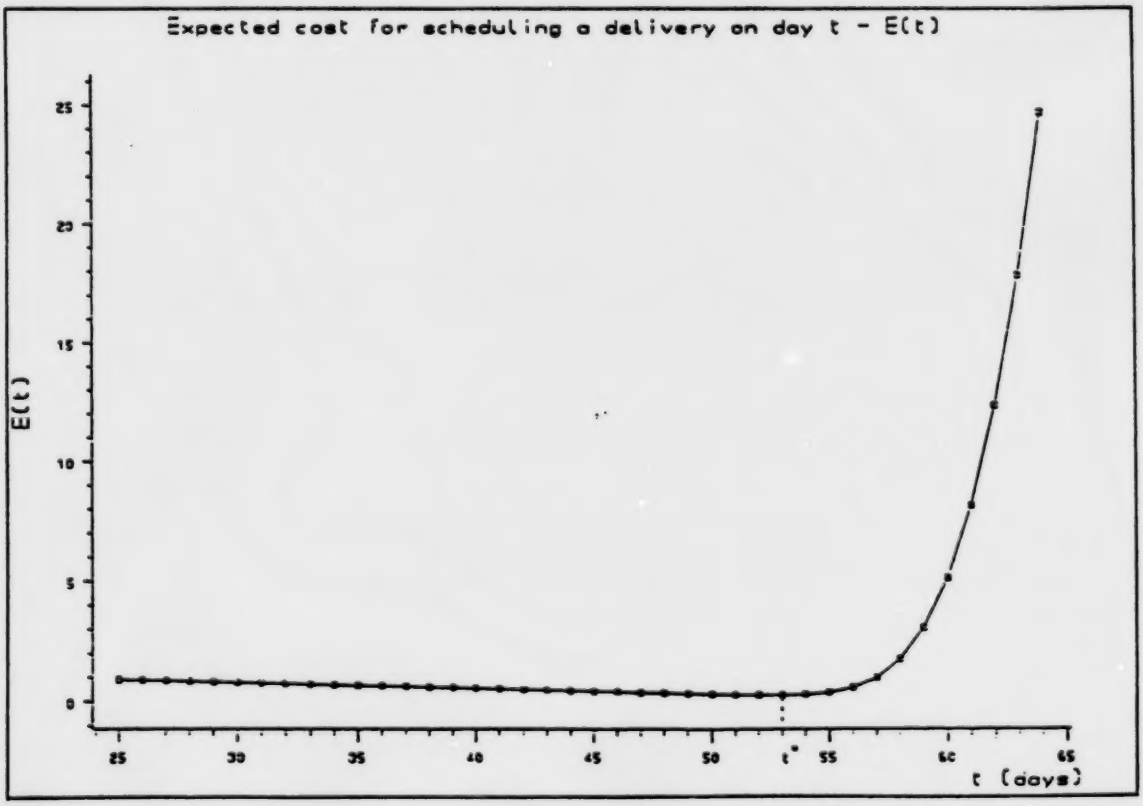

Figure 2 : A typical E(t) function.

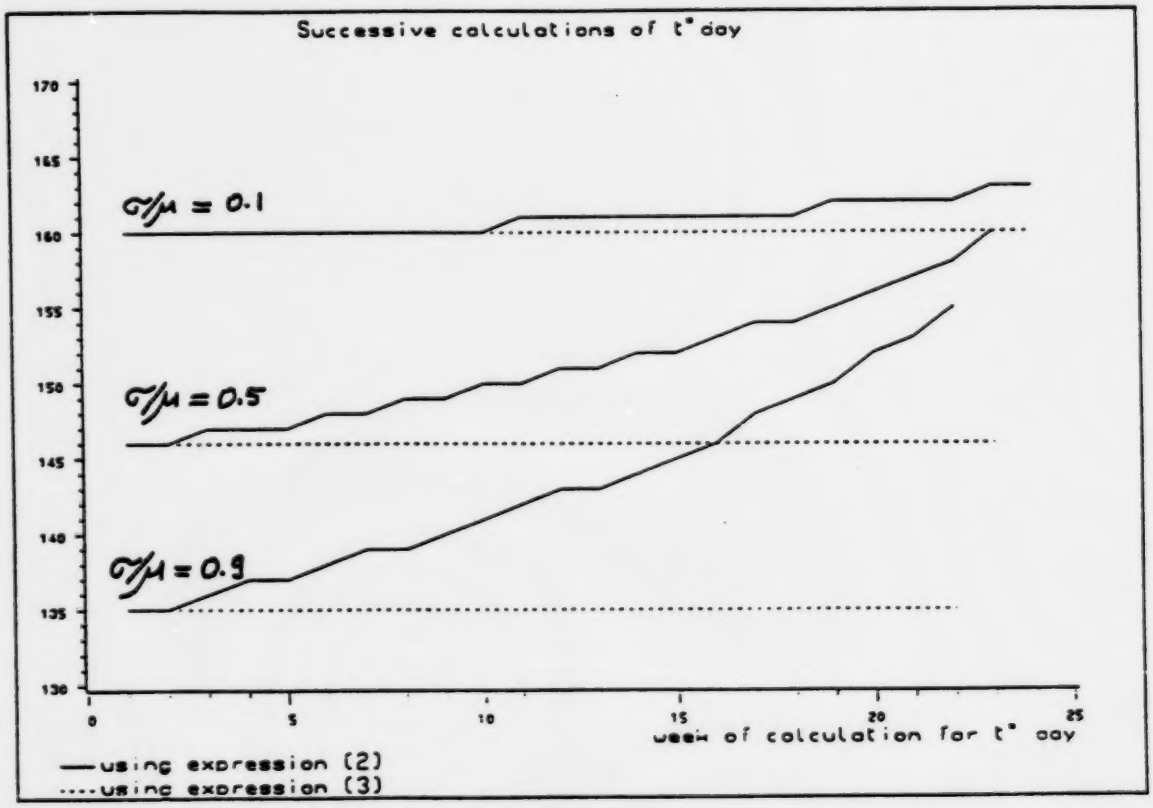

Figure 3 : Successive calculations of t*day.